



US009223654B2

(12) **United States Patent**
Molaro et al.

(10) **Patent No.:** **US 9,223,654 B2**
(45) **Date of Patent:** ***Dec. 29, 2015**

(54) **RESILIENT DISTRIBUTED REPLICATED
DATA STORAGE SYSTEM**

(71) Applicant: **DataDirect Networks, Inc.**, Chatsworth,
CA (US)

(72) Inventors: **Donald J. Molaro**, Cupertino, CA (US);
David Fellingner, Westlake Village, CA
(US); **Adam Fried-Gintis**, Lake Balboa,
CA (US); **Dan Olster**, Woodland Hills,
CA (US)

(73) Assignee: **DataDirect Networks, Inc.**, Chatsworth,
CA (US)

(*) Notice: Subject to any disclaimer, the term of this
patent is extended or adjusted under 35
U.S.C. 154(b) by 0 days.

This patent is subject to a terminal dis-
claimer.

(21) Appl. No.: **14/482,441**

(22) Filed: **Sep. 10, 2014**

(65) **Prior Publication Data**

US 2014/0380093 A1 Dec. 25, 2014

Related U.S. Application Data

(63) Continuation of application No. 14/136,254, filed on
Dec. 20, 2013, now Pat. No. 8,843,447, which is a
continuation-in-part of application No. 13/715,519,
filed on Dec. 14, 2012, now Pat. No. 9,122,699.

(51) **Int. Cl.**

G06F 17/30 (2006.01)

G06F 11/10 (2006.01)

G06F 3/06 (2006.01)

(52) **U.S. Cl.**

CPC **G06F 11/1088** (2013.01); **G06F 3/065**
(2013.01); **G06F 3/067** (2013.01);

(Continued)

(58) **Field of Classification Search**

CPC G06F 11/1088; G06F 11/1076; G06F
11/108; G06F 11/1096; G06F 3/0617; G06F
3/0619; G06F 3/065; G06F 3/0665; G06F
3/0637

USPC 707/640, 652, 690, 691, 821, 827

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,092,732 A 5/1978 Ouchi
4,761,785 A 8/1988 Clark et al.

(Continued)

OTHER PUBLICATIONS

le.;3qRodrigues et al., "High Availability in DHTs: Erasure Coding
vs. Replication", Peer-to-Peer Systems IV 4th International Work-
shop IPTPS 2005 Published Feb. 2005, accessed May 9, 2013, [http://
pmg.csail.mit.edu/papers/rodrigo_IPTPS05.pdf](http://pmg.csail.mit.edu/papers/rodrigo_IPTPS05.pdf).

(Continued)

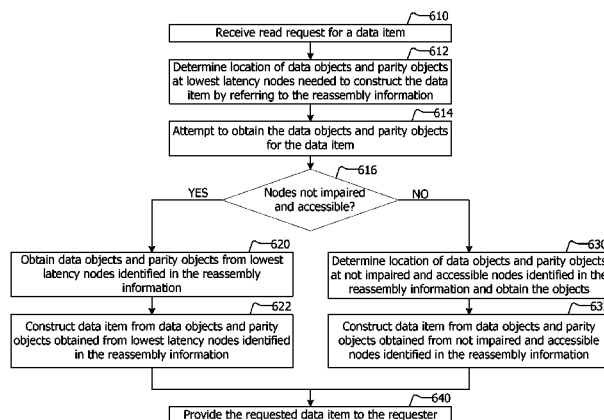
Primary Examiner — Hares Jami

(74) *Attorney, Agent, or Firm* — SoCal IP Law Group LLP;
Mark A. Goldstein

(57) **ABSTRACT**

A resilient distributed replicated data storage system is
described herein. The storage system includes zones that are
independent, and autonomous from each other. The zones
include nodes that are independent and autonomous. The
nodes include storage devices. When a data item is stored, it
is partitioned into a plurality of data objects and a plurality of
parity objects are calculated. Reassembly instructions are
created for the data item. The data objects, parity objects and
reassembly instructions are spread across nodes and zones in
the storage system according to a policy for the data item.
When a zone is inaccessible, a virtual zone is created and used
until the intended zone is available. When a read request is
received, the data item is prepared from the lowest latency
nodes according to the reassembly instructions, and a virtual
zone is accessed in place of a real zone when the real zone is
inaccessible.

22 Claims, 10 Drawing Sheets



(52) U.S. Cl.

CPC G06F 3/0617 (2013.01); G06F 3/0619 (2013.01); G06F 3/0665 (2013.01); G06F 11/108 (2013.01); G06F 11/1076 (2013.01); G06F 11/1096 (2013.01); G06F 3/0637 (2013.01)

(56)

References Cited

U.S. PATENT DOCUMENTS

5,208,813	A	5/1993	Stallmo	
5,390,187	A	2/1995	Stallmo	
5,504,892	A	4/1996	Atsatt et al.	
5,758,153	A	5/1998	Atsatt et al.	
6,154,853	A	11/2000	Keden	
6,442,659	B1	8/2000	Blumenau	
6,658,439	B2	12/2003	Karasudani	
RE39,421	E	12/2006	Stolowitz	
7,529,970	B2	5/2009	Hartline et al.	
7,734,643	B1 *	6/2010	Waterhouse G06F 17/30067 707/770
8,606,958	B1 *	12/2013	Evans A63F 13/12 709/203
2001/0016841	A1	8/2001	Karasudani	
2002/0178162	A1	11/2002	Ulrich et al.	
2003/0115438	A1	6/2003	Mahalingam et al.	
2004/0098383	A1	5/2004	Tabellion et al.	
2004/0107314	A1	6/2004	Kim et al.	
2004/0236769	A1	11/2004	Smith et al.	
2005/0216502	A1	9/2005	Kaura et al.	
2009/0106255	A1	4/2009	Lacpra et al.	
2009/0198719	A1	8/2009	DeWitt	
2010/0180153	A1	7/2010	Jernigan et al.	
2010/0218037	A1 *	8/2010	Swartz G06F 17/30 714/6.12
2010/0223495	A1	9/2010	Leppard	
2011/0153977	A1	6/2011	Root et al.	
2011/0196900	A1 *	8/2011	Drobychev G06F 17/30082 707/812
2011/0214011	A1 *	9/2011	Grube G06F 11/1076 714/6.22
2011/0219048	A1	9/2011	Cross et al.	
2011/0276622	A1	11/2011	Knapp et al.	
2012/0030736	A1	2/2012	Resch et al.	
2013/0346532	A1 *	12/2013	D'Amato G06F 3/0617 709/213

OTHER PUBLICATIONS

Huang et al., "Erasure Coding in Windows Azure Storage", Microsoft Corporation, 2012 Usenic best paper, accessed May 9, 2013, http://research.microsoft.com/en-us/um/people/yekhanin/Papers/UsenixATC_2012.pdf.

Gsoedl, "Inside object-based storage", SearchStorage.com, article published Jun. 2012, accessed Oct. 5, 2012, <http://searchstorage.techtarget.com/magazineContent/Inside-object-based-storage?vgnextfmt=print>, pp. 1-3.

Gsoedl, "Advantages of using an object storage system", SearchStorage.com, article published Jul. 30, 2012, accessed Oct. 5, 2012, <http://searchcloudstorage.techtarget.com/tip/Advantages-of-using-an-object-storage-system?vgnextfmt=print>, pp. 1-3.

Primmer, "Distributed Object Store Principles of Operation: The Case for Intelligent Storage", White Paper, Published Jul. 2010 by Hitachi Data Systems, total of 20 pages.

Terrace et al., "Object Storage on CRAQ, High-throughput chain replication for read-mostly workloads", In Proc. USENIX Annual Technical Conference, San Diego, CA, Jun. 2009, total of 16 pages.

Taylor, "Object-based storage promises scalability and flexibility for static, unstructured data needs", SearchStorage.co.UK, Apr. 11, 2011, total of 2 pages.

Slack, "Object-based storage for cloud storage: Product assessments", searchCloudProvider.com, published Oct. 26, 2010, 2 total pages.

Slack, "Object storage: Object-based storage devices challenge file systems for unstructured data sets", SearchStorageChannel.com, published Sep. 22, 2010, 2 total pages.

Carns et al., "Object storage semantics for replicated concurrent-writer file systems", journal, Argonne National Laboratory, accessed Oct. 5, 2012, http://www.mcs.anl.gov/events/workshops/iasds10/carns_iasds2010.pdf.

Factor et al., "Object Storage: The Future Building Block for Storage Systems", A Position Paper, IBM Haifa Research Laboratories, pp. 101-105, accessed Oct. 5, 2012, <https://www.research.ibm.com/haifa/projects/storage/objectstore/papers/PositionOSD.pdf>.

Rackspace, Inc., "CloudFiles Introduction", Jan. 24, 2011, 13 total pages.

DataDirect Networks, Inc., "Web Object Scaler", Brochure, Mar. 9, 2010, total of 4 pages.

Connor et al., "Cloud Storage: Adoption, Practice and Deployment", Storage Strategies NOW, Outlook Report, Published Apr. 4, 2011, total of 112 pages.

Berriman, et al., Netapp Raid-DP: Dual-Parity Raid 6 Protection Without Compromise, WP-7005-1006, Oct. 2006, pp. 1-11.

Celeros, Bringing SANity to Storage Costs, RAID White Paper, Jan. 2007.

Chen et al., RAID: High-Performance, Reliable Secondary Storage, AMC Computing Surveys, vol. 26, No. 2, Jun. 1994, pp. 145-185.

DataDirect Networks, DataDirect Networks S2A9550: High-Performance, Resilient Storage for High-Performance Computing (HPC), White Paper, Sep. 2007, pp. 1-12.

DataDirect Networks, Best Practices: Enterprise SATA Deployment with High Performance and Reliability, White Paper, Dec. 2007, pp. 1-12.

EMC, EMC CLARiiON Raid 6 Technology, White Paper, Jul. 2007, pp. 1-14.

MPSTOR, RAID 50 Interleave, White Paper, Xtore Extreme Storage, Nov. 2006, pp. 1-13.

Patterson, et al., A Case for Redundant Arrays of Inexpensive Disks (RAID), Computer Science Division, Department of Electrical Engineering and Computer Sciences, pp. 1-25.

Pivot3, Pivot3 RAIGE Storage Cluster, White Paper, Technology Overview, Feb. 2007, pp. 1-18.

STORAGE, Best Storage Products of 2007, issued Feb. 2008, accessed on Apr. 22, 2008, http://searchstorage.techtarget.com/magazinePrintFriendly/0,296905,sid5_gci1299110,00.html.

WinchesterSystems, Technology Update White Paper "Enterprise RAID 6", May 4, 2006, pp. 1-14.

Du, et al., Experiences Building and Object-Based Storage System Based on the OSD T-10 Standard, DTC Intelligent Storage Consortium, University of Minnesota, 2006, pp. 1-11.

Moore, et al., Storage Resource Broker Global Data Grids, San Diego Supercomputer Center, 2006, pp. 1-12.

Devulapalli, et al., Integrating Parallel File Systems with Object-Based Storage Devices, Ohio Supercomputer Center, pp. 1-10.

Smolik, An Object-Oriented File System—An Example of Using the Class Hierarchy Framework Concept, Department of Computer Sciences, University of Missouri—Rolla, pp. 33-53.

Olson, et al., Secure Capabilities for a Petabyte-Scale object-Based Distributed File System, Storage System Research Center, Computer Science Department, University of California, Santa Cruz, pp. 64-73.

Marsden et al., Improving the Usability of the Hierarchical File System, Department of Computer Science, University of Cape Town and University of Stirling, Proceedings of SAICSIT 2003, pp. 122-129.

Gaffey, SGI's Cluster File System—CXFS, File Systems Engineering, Apr. 9, 2000, slides 1-56.

Foster, Global data Services, Developing Data-Intensive Applications Using Globus Software, Computation Institute Argonne National Lab & University of Chicago, slides 1-114.

Pollack et al., Efficient Access Control for Distributed Hierarchical File Systems, University of California, Santa Cruz, slides 1-9.

Pollack et al., Efficient Access Control for Distributed Hierarchical File Systems, University of California, Santa Cruz, 2005, Journal, pp. 101-108.

Shinkai et al., Alternative Implementations of Cluster File Systems, MMS Conference, Mar. 2000, slides 1-16.

(56)

References Cited

OTHER PUBLICATIONS

Rogers et al., z/OS Distributed File Service zSeries File System Implementation z/OS V1R11, Oct. 2009, Redbooks, IBM, International Technical Support Organization, Draft Document, Nov. 24, 2009, pp. 1-458.

Pike et al., The Styx Architecture for Distributed Systems, Computing Science Research Center, Bell Labs Technical Journal, vol. 4, No. 2, Apr.-Jun. 1999, pp. 1-11.

Welch et al., Object-Based Cluster Storage Systems, Panasas, Inc., May 25, 2006, slides 1-104.

* cited by examiner




FIG. 1

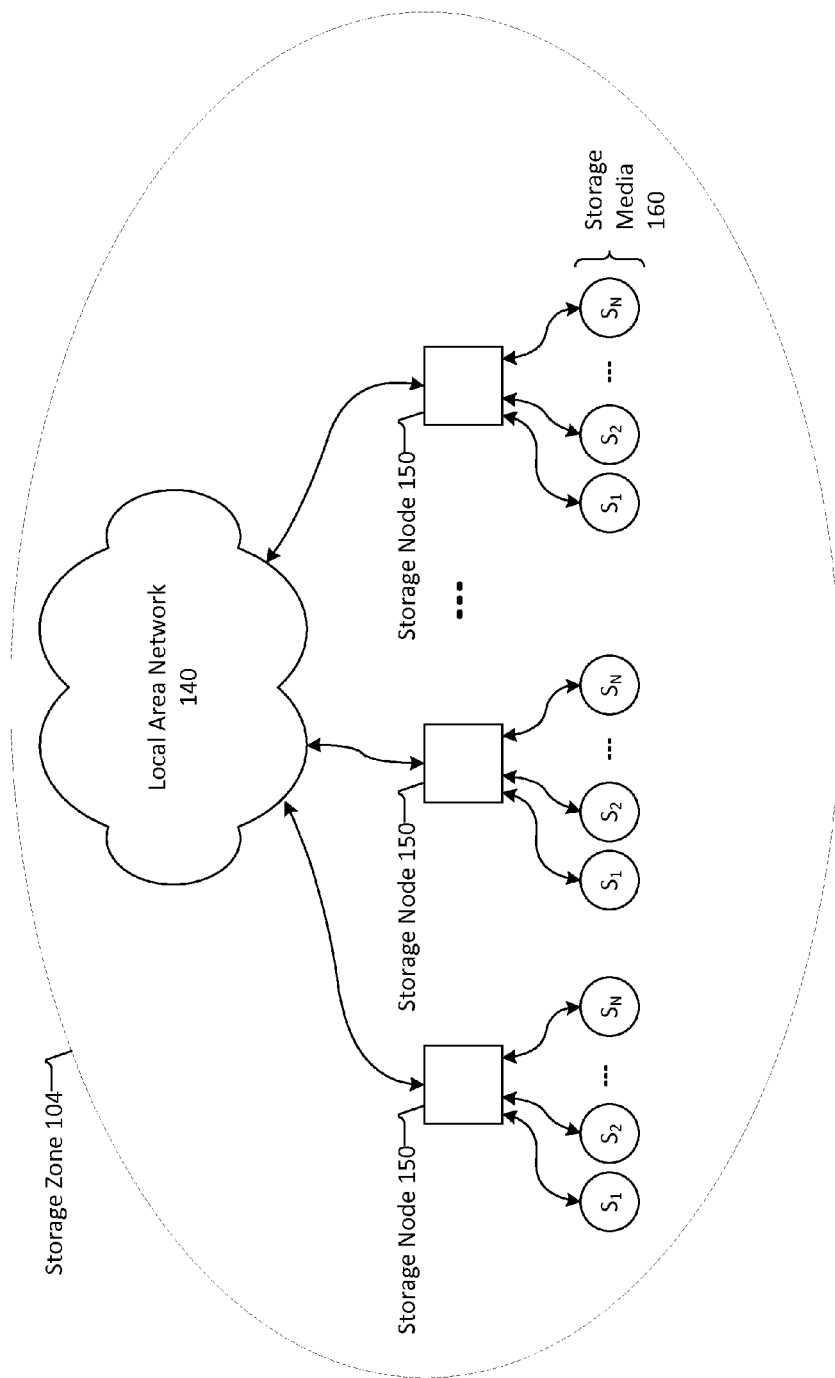


FIG. 2

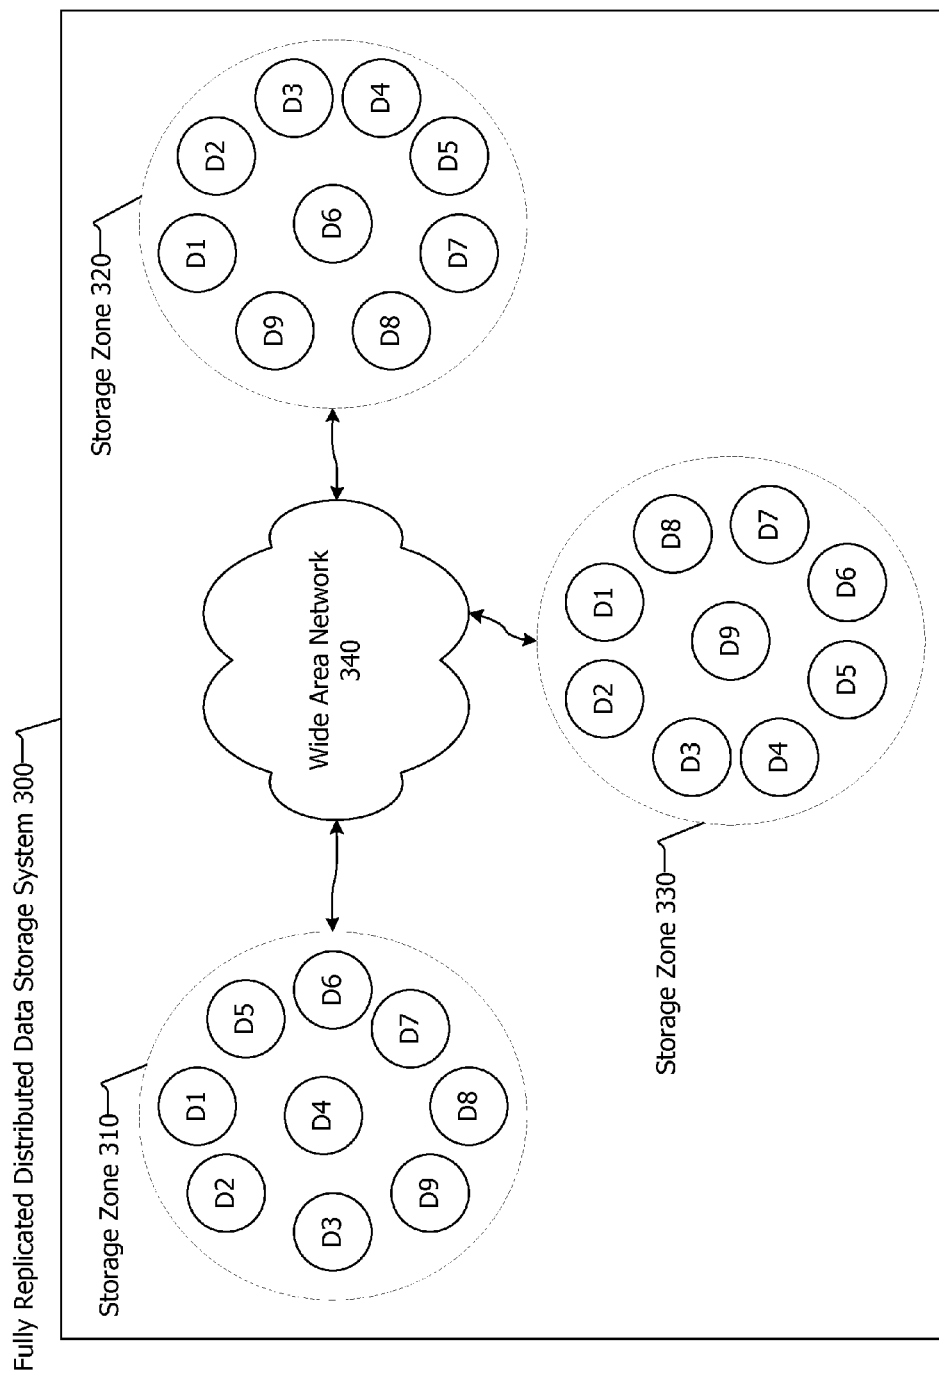


FIG. 3

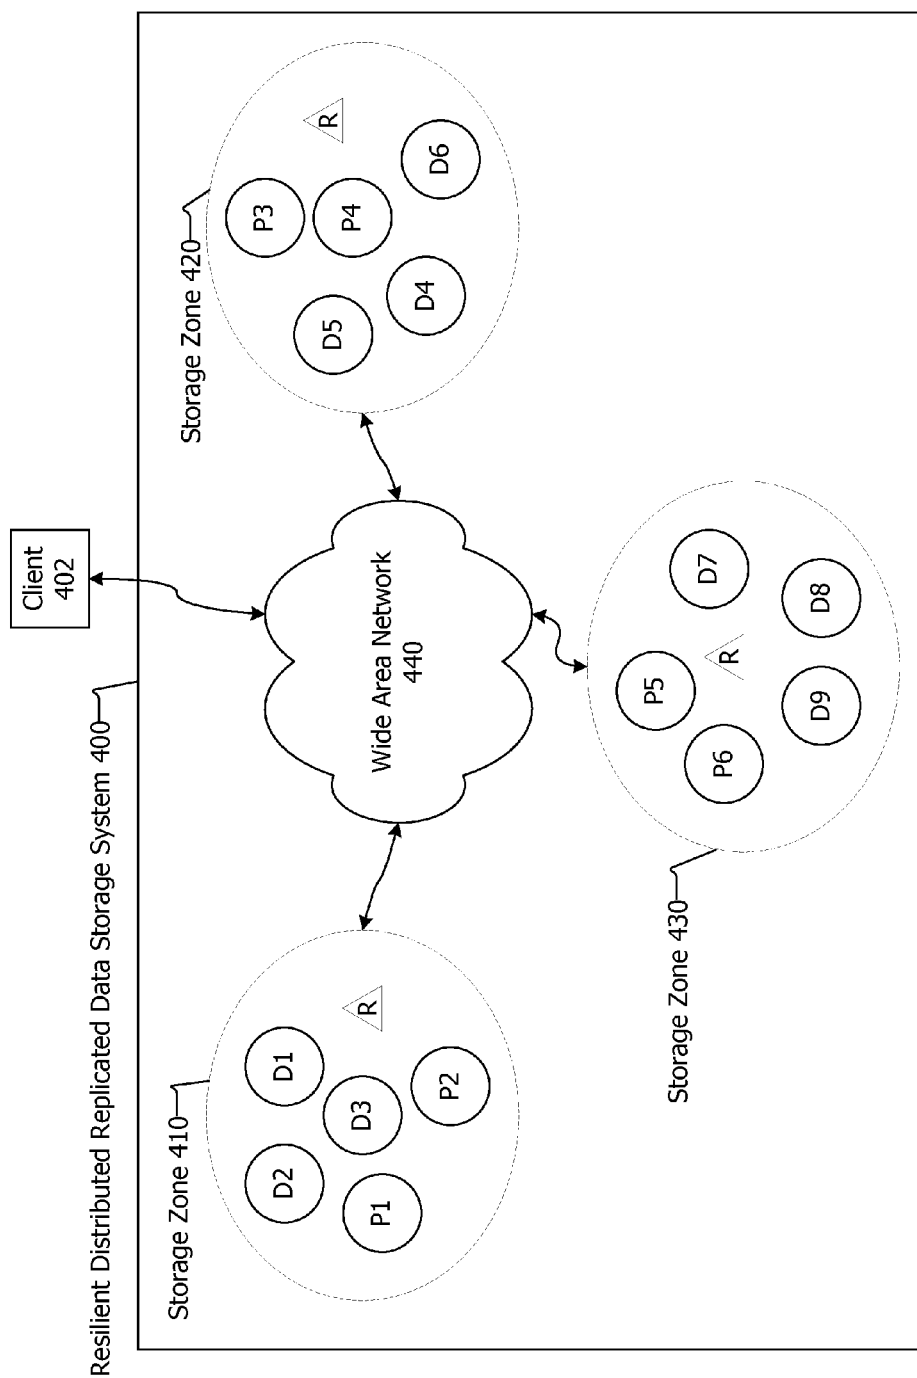


FIG. 4

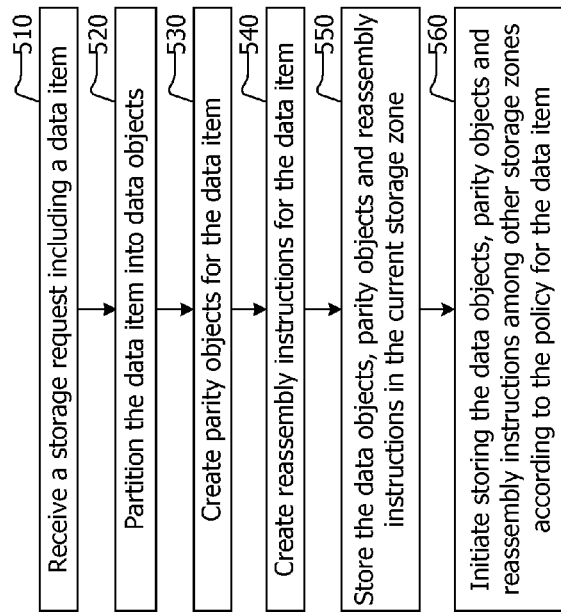


FIG. 5

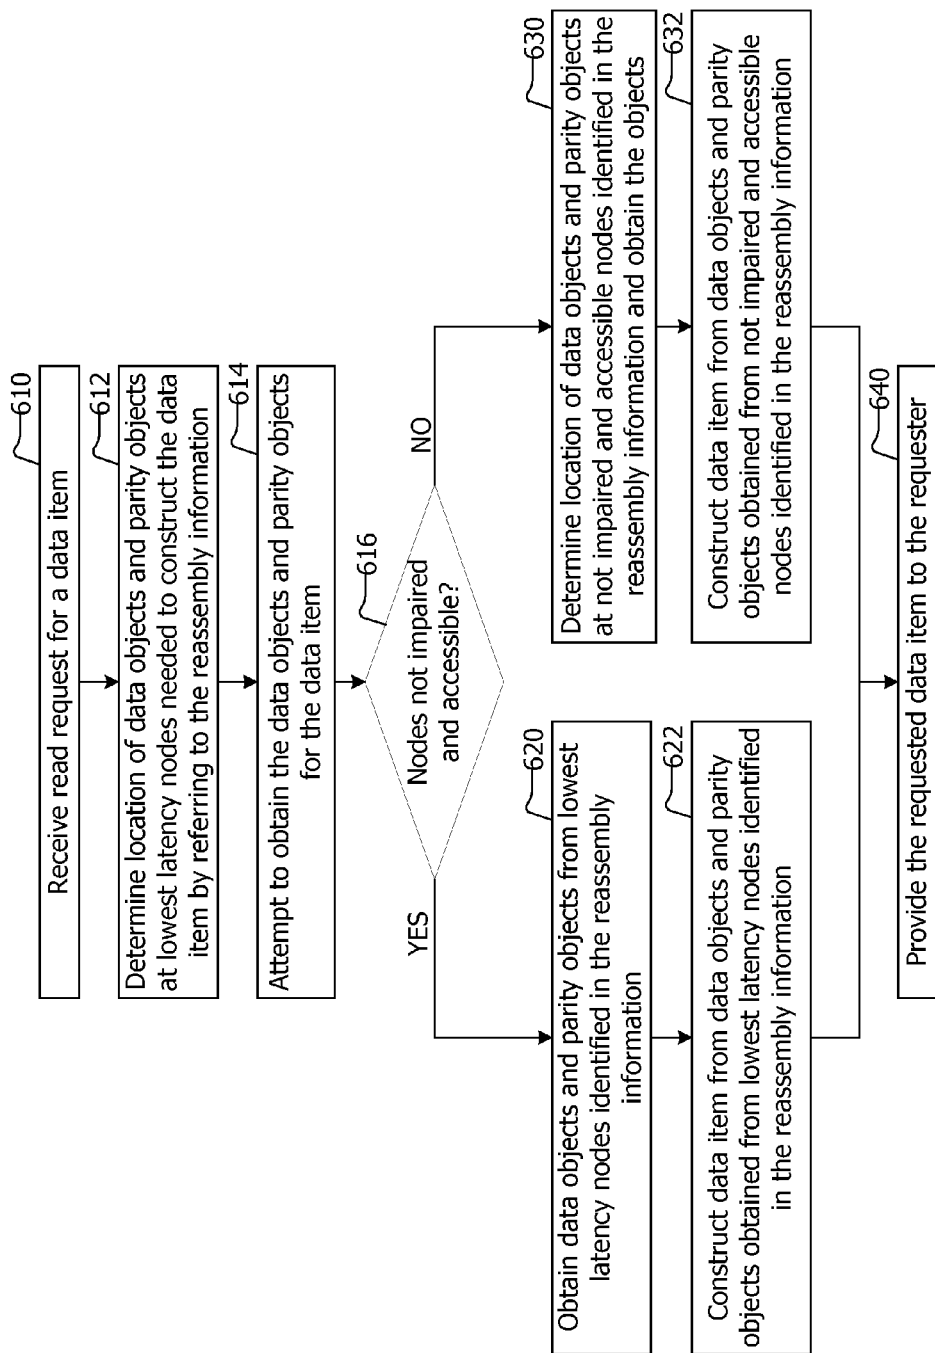


FIG. 6

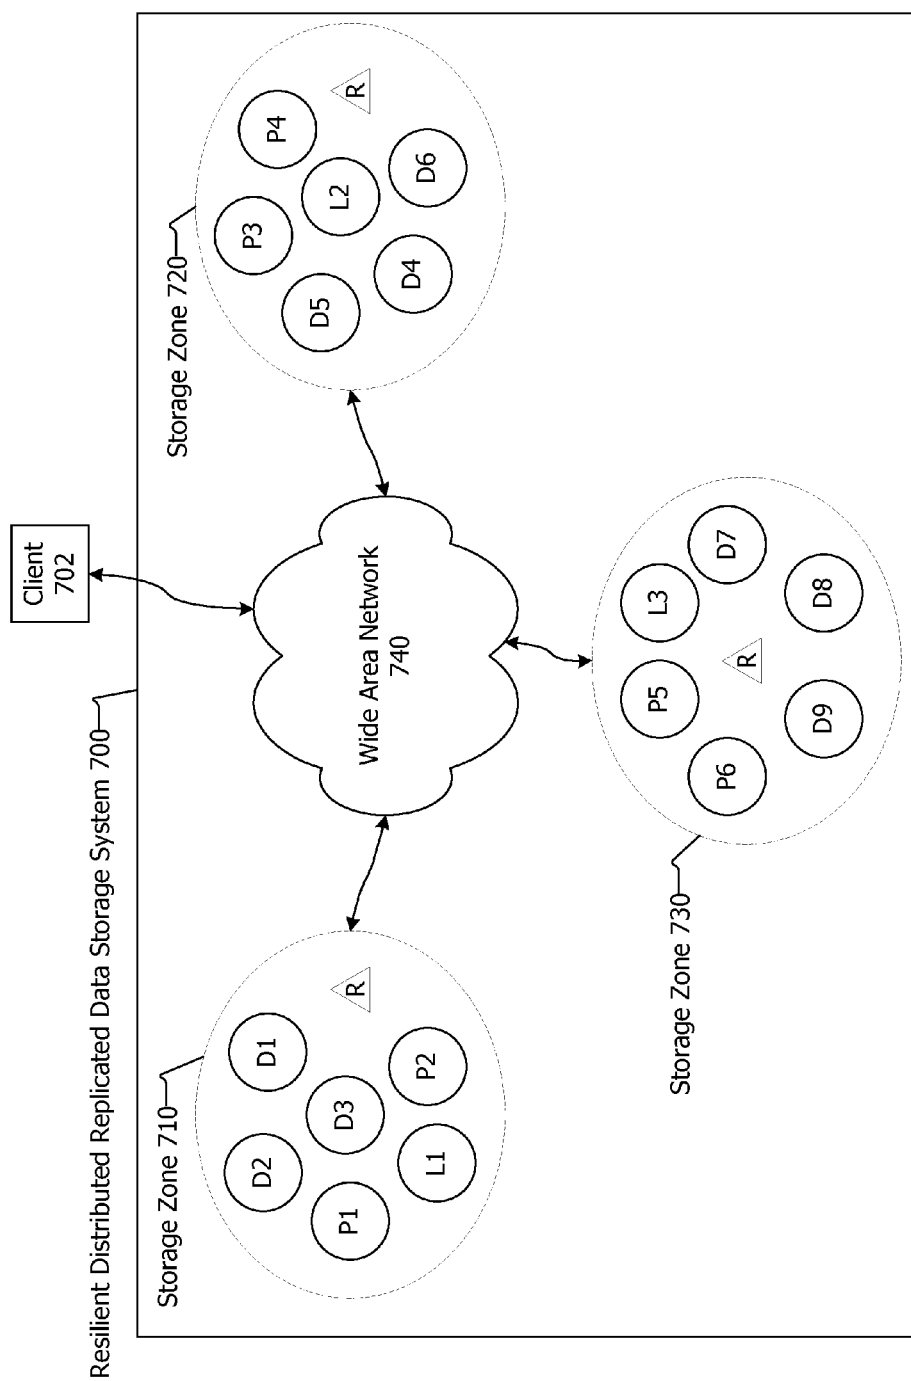


FIG. 7

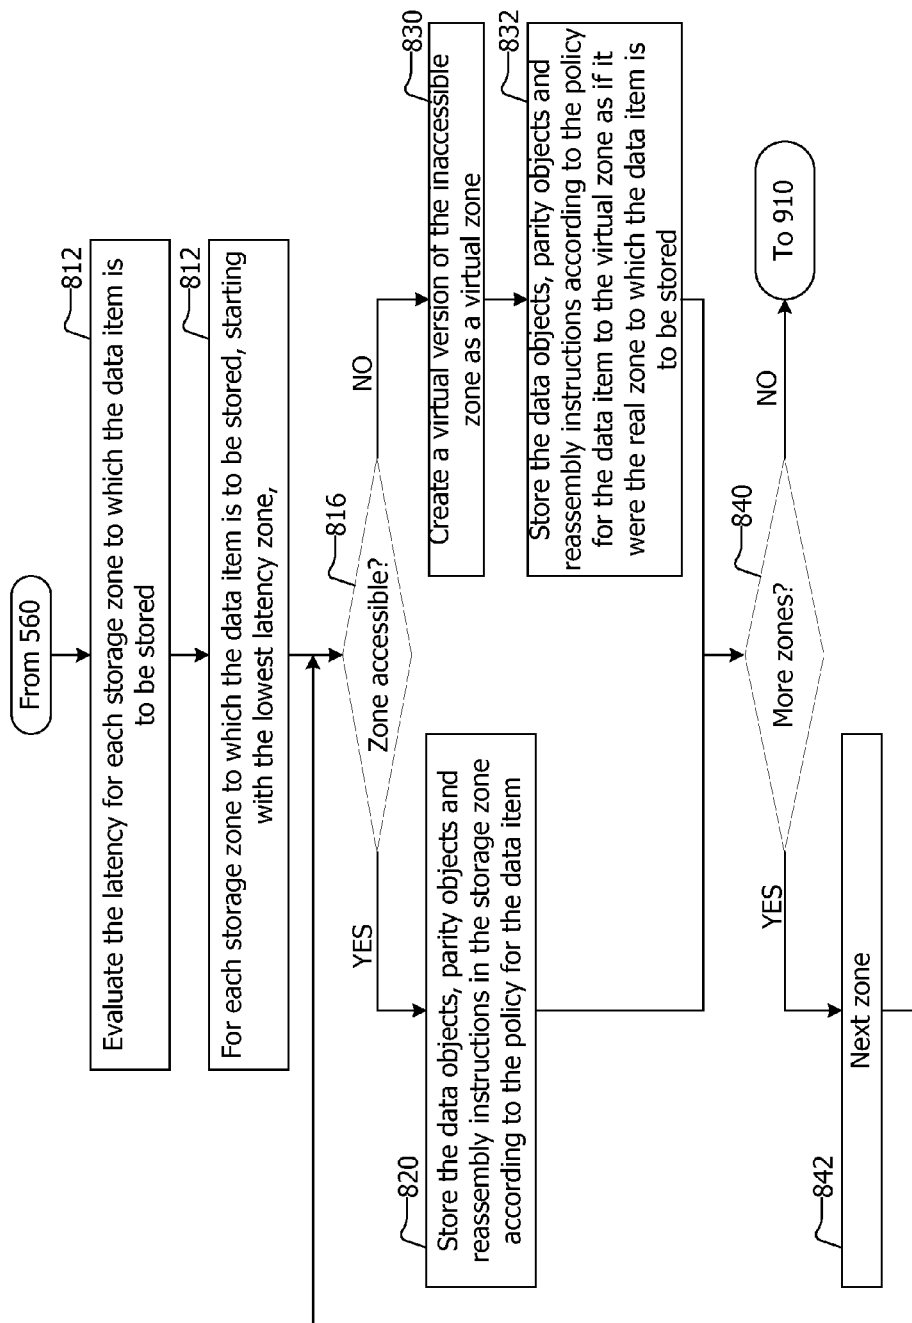


FIG. 8




FIG. 9

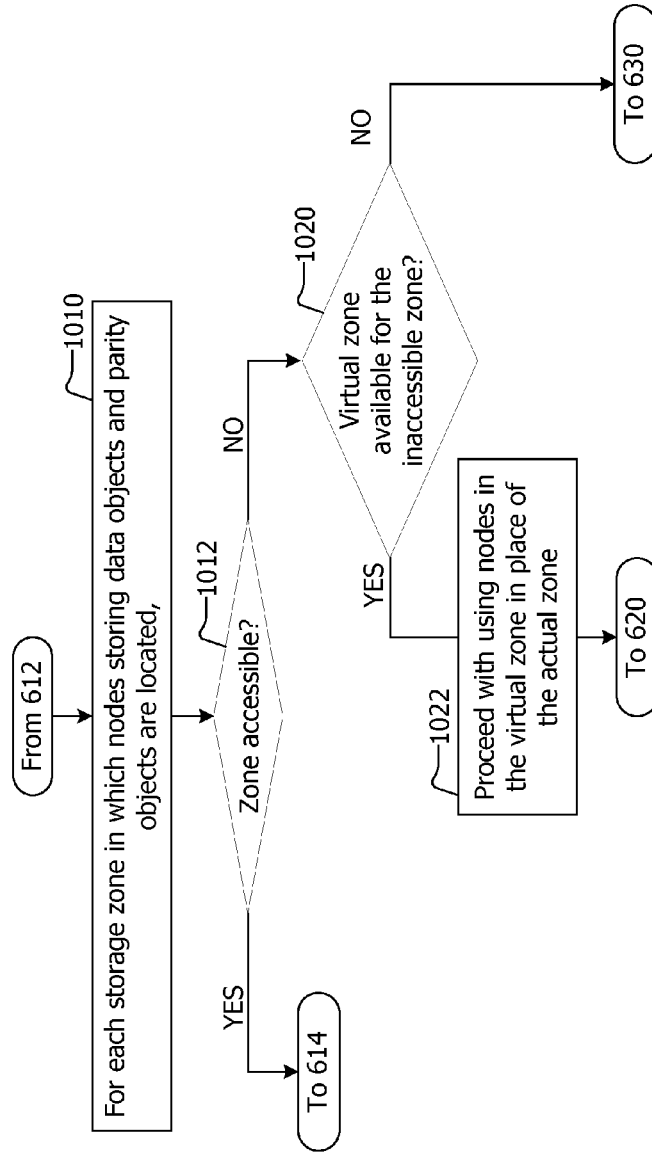


FIG. 10

1

RESILIENT DISTRIBUTED REPLICATED DATA STORAGE SYSTEM

RELATED APPLICATION INFORMATION

This patent is a continuation of U.S. application Ser. No. 14/136,254 filed Dec. 20, 2013 entitled RESILIENT DISTRIBUTED REPLICATED DATA STORAGE SYSTEM, now U.S. Pat. No. 8,843,447 issued Sep. 23, 2014, which is a continuation in part of pending U.S. application Ser. No. 13/715,519 filed Dec. 14, 2012 entitled FAILURE RESILIENT DISTRIBUTED REPLICATED DATA STORAGE SYSTEM.

NOTICE OF COPYRIGHTS AND TRADE DRESS

A portion of the disclosure of this patent document contains material which is subject to copyright protection. This patent document may show and/or describe matter which is or may become trade dress of the owner. The copyright and trade dress owner has no objection to the facsimile reproduction by anyone of the patent disclosure as it appears in the Patent and Trademark Office patent files or records, but otherwise reserves all copyright and trade dress rights whatsoever.

BACKGROUND

1. Field

This disclosure relates to data stored in a distributed replicated data storage system and an improved resilient method for storing and accessing data in a distributed replicated data storage system.

2. Description of the Related Art

A file system is used to store and organize computer data stored as electronic files. File systems allow files to be found, read, deleted, and otherwise accessed. File systems store files on one or more storage devices. File systems store files on storage media such as hard disk drives and silicon storage devices.

Various applications may store large numbers of documents, images, audio, videos and other data as objects using a distributed replicated data storage system in which data is replicated and stored in at least two locations.

DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of a distributed replicated data storage system.

FIG. 2 is a block diagram of a storage zone included in a distributed replicated data storage system.

FIG. 3 is a block diagram of an example fully replicated distributed data storage system.

FIG. 4 is a block diagram of a first example resilient distributed replicated data storage system.

FIG. 5 is a flow chart of the actions taken to store or put data in a resilient distributed replicated data storage system.

FIG. 6 is a flow chart of the actions taken to read or get data included in a resilient distributed replicated data storage system.

FIG. 7 is a block diagram of a second example resilient distributed replicated data storage system.

FIG. 8 is a flow chart of additional actions taken to store or put data in a resilient distributed replicated data storage system.

FIG. 9 is a flow chart of the further additional actions taken to store or put data in a resilient distributed replicated data storage system.

2

FIG. 10 is a flow chart of additional actions taken to read or get data included in a resilient distributed replicated data storage system.

DETAILED DESCRIPTION

Environment

FIG. 1 is a block diagram of a distributed replicated data storage system **100**. The distributed replicated data storage system **100** includes at least two storage zones. The distributed replicated data storage system **100** includes multiple storage zones that are independent of one another, may be geographically dispersed and are in a peer-to-peer configuration. In the example shown, the distributed replicated data storage system **100** includes two storage zones, first storage zone **110** and second storage zone **120**.

The storage zones **110** and **120** are separated geographically. The storage zones **110** and **120** communicate with each other and share objects over wide area network **130**. The wide area network **130** may be or include the Internet. The wide area network **130** may be wired, wireless, or a combination of these. The wide area network **130** may be public or private, may be a segregated network, and may be a combination of these. The wide area network **130** includes networking devices such as routers, hubs, switches and the like.

The term data as used herein includes a bit, byte, word, block, stripe or other unit of information. In one embodiment the data is stored within and by the distributed replicated data storage system as objects. As used herein, the term data is inclusive of entire computer readable files or portions of a computer readable file. The computer readable file may include or represent text, numbers, data, images, photographs, graphics, audio, video, computer programs, computer source code, computer object code, executable computer code, and/or a combination of these and similar information.

Many data intensive applications store a large quantity of data, these applications include scientific applications, newspaper and magazine websites (for example, nytimes.com and life.com), scientific lab data capturing and analysis programs, video and film creation software, and consumer web based applications such as social networking websites (for example, FACEBOOK), photo sharing websites (for example, FLIKR), video sharing websites (for example, YOUTUBE) and music distribution websites (for example, ITUNES).

FIG. 2 is a block diagram of a storage zone **104** included in a distributed replicated data storage system. The storage zones **110** and **120** shown in FIG. 1 are examples of storage zone **104**. The storage nodes **150** within a storage zone **104** may be connected via a local area network **140** by wire lines, optical fiber cables, wireless communication connections, and others, and may be a combination of these. The local area network **140** may include one or more networking devices such as routers, hubs, switches and the like.

Referring again to FIG. 1, one node in one of the storage zones may include an object management system and/or application program that automatically identifies and adapts to impairments in the storage nodes **150**. That is, software in a node in the storage zone **110** may monitor and recognize multiple kinds of impairments, including, for example, storage node failures; storage device failures; storage nodes with diminished capacity such as one or more down or diminished drives; storage devices with diminished capacity such as bad blocks or sectors; other equipment failures; offline nodes; offline storage devices; missing data; impaired data; and other system anomalies. When the software in a node in the storage zone recognizes an impairment, the distributed replicated

data storage system may automatically adapt. In a related embodiment, the distributed replicated data storage system **100** may include an application program that automatically identifies impairments in constituent storage zones and takes remedial action.

The storage zones **110**, **120** and **104** may include a computing device and/or a controller on which software may execute. The computing device and/or controller may include one or more of logic arrays, memories, analog circuits, digital circuits, software, firmware, and processors such as microprocessors, field programmable gate arrays (FPGAs), application specific integrated circuits (ASICs), programmable logic device (PLDs) and programmable logic array (PLAs). The hardware and firmware components of the computing device and/or controller may include various specialized units, circuits, software and interfaces for providing the functionality and features described herein. The processes, functionality and features described herein may be embodied in whole or in part in software which operates on a controller and/or one or more computing devices and may be in the form of one or more of firmware, an application program, object code, machine code, an executable file, an applet, a COM object, a dynamic linked library (DLL), a script, one or more subroutines, or an operating system component or service, and other forms of software. The hardware and software and their functions may be distributed such that some actions are performed by a controller or computing device, and others by other controllers or computing devices within a storage zone.

A computing device as used herein refers to any device with a processor, memory and a storage device that may execute instructions such as software including, but not limited to, server computers, personal computers, portable computers, and laptop computers. The computing devices may run an operating system, including, for example, versions of the Linux, Unix, MS-DOS, Microsoft Windows, Solaris, Symbian, Android, Chrome, and Apple Mac OS X operating systems. Computing devices may include a network interface in the form of a card, chip or chip set that allows for communication over a wired and/or wireless network. The network interface may allow for communications according to various protocols and standards, including, for example, versions of Ethernet, INFINIBAND® network, Fibre Channel, and others. A computing device with a network interface is considered network capable.

Referring again to FIG. 2, the storage zone **104** includes a plurality of storage nodes **150** which include a plurality of storage media **160**. Each of the storage nodes **150** may include one or more server computers. Each of the storage nodes **150** may be an independent network attached storage (NAS) device or system. The term "storage media" is used herein to refer to any configuration of hard disk drives, solid-states drives, silicon storage devices, magnetic tape, or other similar non-transitory storage media. Hard disk drives and/or other storage media **160** may also be arranged in the storage nodes **150** according to any of a variety of techniques.

The storage media included in a storage node may be of the same capacity, may have the same physical size, and may conform to the same specification, such as, for example, a hard disk drive specification. Example sizes of storage media include, but are not limited to, 2.5" and 3.5". Example hard disk drive capacities include, but are not limited to, 500 Mbytes, 1 terabyte and 2 terabytes. Example hard disk drive specifications include Serial Attached Small Computer System Interface (SAS), Serial Advanced Technology Attachment (SATA), and others. An example storage node may include 16 one terabyte 3.5" hard disk drives conforming to the SATA standard. In other configurations, the storage nodes

150 may include more and fewer drives, such as, for example, 10, 12, 24, 32, 40, 48, 64, etc. In other configurations, the storage media **160** in a storage node **150** may be hard disk drives, silicon storage devices, magnetic tape devices, or a combination of these. In some embodiments, the physical size of the media in a storage node may differ, and/or the hard disk drive or other storage specification of the media in a storage node may not be uniform among all of the storage devices in a storage node **150**.

The storage media **160** in a storage node **150** may be included in a single cabinet, rack, shelf or blade. When the storage media in a storage node are included in a single cabinet, rack, shelf or blade, they may be coupled with a backplane. A controller may be included in the cabinet, rack, shelf or blade with the storage devices. The backplane may be coupled with or include the controller. The controller may communicate with and allow for communications with the storage media according to a storage media specification, such as, for example, a hard disk drive specification. The controller may include a processor, volatile memory and non-volatile memory. The controller may be a single computer chip such as an FPGA, ASIC, PLD and PLA. The controller may include or be coupled with a network interface.

In another embodiment, multiple storage nodes **150** are included in a single cabinet or rack such that a storage zone may be included in a single cabinet. When in a single cabinet or rack, storage nodes and/or constituent storage media may be coupled with a backplane. A controller may be included in the cabinet with the storage media and/or storage nodes. The backplane may be coupled with the controller. The controller may communicate with and allow for communications with the storage media. The controller may include a processor, volatile memory and non-volatile memory. The controller may be a single computer chip such as an FPGA, ASIC, PLD and PLA.

The rack, shelf or cabinet containing a storage zone may include a communications interface that allows for connection to other storage zones, a computing device and/or to a network. The rack, shelf or cabinet containing a storage node **150** may include a communications interface that allows for connection to other storage nodes, a computing device and/or to a network. The communications interface may allow for the transmission of and receipt of information according to one or more of a variety of standards, including, but not limited to, universal serial bus (USB), IEEE 1394 (also known as FIREWIRE® and i.LINK®), Fibre Channel, Ethernet, WiFi (also known as IEEE 802.11). The backplane or controller in a rack or cabinet containing a storage zone may include a network interface chip, chipset, card or device that allows for communication over a wired and/or wireless network, including Ethernet. The backplane or controller in a rack or cabinet containing one or more storage nodes **150** may include a network interface chip, chipset, card or device that allows for communication over a wired and/or wireless network, including Ethernet. In various embodiments, the storage zone, the storage node, the controller and/or the backplane may provide for and support 1, 2, 4, 8, 12, 16, etc. network connections and may have an equal number of network interfaces to achieve this.

The techniques discussed herein are described with regard to storage media including, but not limited to, hard disk drives and solid-state drives. The techniques may be implemented with other readable and writable storage media.

As used herein, a storage device is a device that allows for reading from and/or writing to a storage medium. Storage devices include hard disk drives (HDDs), solid-state drives (SSDs), DVD drives, flash memory devices, and others. Stor-

age media include magnetic media such as hard disks and tape, flash memory, and optical disks such as CDs, DVDs and BLU-RAY® discs.

In some embodiments, files and other data may be partitioned into smaller portions and stored as multiple objects among multiple storage media **160** in a storage node **150**. Files and other data may be partitioned into smaller portions referred to as objects and stored among multiple storage nodes **150** in a storage zone. In one embodiment each object includes a storage policy identifier and a data portion. The object including its constituent data portion is stored among storage nodes and storage zones according to the storage policy specified by the storage policy identifier included in the object. Various policies may be maintained and distributed or known to the nodes in all zones in the distributed replicated storage system. Policies define the replication and placement of data objects in the data storage system. Example policies include, full distribution, single copy, single copy to a specific zone, copy to all zone but a specified zone, and others. The policies may be stored on and distributed from a client **102** to the distributed replicated storage system **100** to all nodes in the storage system. A character (e.g., A, B, C, etc.) or number (0, 1, 2, etc.) or combination of one or more characters and numbers (A1, AAA, A2, BC3, etc.) or other scheme may be associated with and used to identify each of the policies. The client **102** of the storage system **100** may be a computing device such as, for example, a personal computer, tablet, mobile phone, workstation or server.

Referring again to FIG. 1, the wide area network **130** connects geographically separated storage zones. Each of the storage zones includes a local area network **140**. The transfer of data between storage nodes in a storage zone is fast compared to communication over the wide area network **130**.

Referring now to FIG. 3, a fully replicated distributed data storage system **300** is shown. In this configuration, all data is replicated among all storage zones such that all copies of stored data are available from and accessible from all storage zones. This is referred to herein as a fully replicated storage system. As shown, a data item is partitioned into data objects D1 through D9. All data objects D1 through D9 are available in all three of storage zones **310**, **320** and **330**. The data item stored as data objects D1 through D9 may represent one movie, one audio file, or any portion, full or partial, of any data. In a fully replicated storage system, all stored data is accessed from any or all of the storage zones based on various system rules. In a fully replicated configuration, the storage zones **310**, **320** and **330** are replicated such that copies of data in all three storage zones are the same. Replication may be performed synchronously, that is, completed before the write operation is acknowledged; asynchronously, that is, the replicas may be written before, after or during the write of the first copy; or a combination of each. Full replication can result in increased storage costs, particularly when there are multiple zones and all data is replicated at all zones. Although this configuration provides for resiliency, it may be costly monetarily and may be too costly monetarily for some. That is, the monetary cost of the constituent storage media required for full replication may be too high.

In the resilient distributed replicated data storage system described herein, when writing data to a storage zone, the data may be replicated in one or more additional storage zones to provide for redundancy such that access to data is possible when a zone goes down or is impaired or unreachable, without the need for full replication. The resilient replication system does not require that each zone have a full copy of all data objects. The techniques described herein result in a lower cost resilient data storage system.

The lower cost is achieved through the use of smaller capacity storage zones (and fewer or smaller capacity storage devices in those storage zones). Because of the monetary cost of data storage devices, the resilient replication system described herein reduces the amount of data stored, requiring a smaller storage capacity which reduces storage media costs.

Referring now to FIG. 4, a resilient distributed replicated storage system **400** is shown. The resilient distributed replicated data storage system **400** includes multiple storage zones that are independent of one another, are geographically dispersed and are in a peer-to-peer configuration. A client computing device **402** may be used to specify the configuration of the distributed replicated storage system **400** and may promulgate policies or storage criteria for the nodes in the system **400**. In this example, a data item is partitioned into nine pieces, data objects D1 through D9. Data objects D1, D2 and D3 are stored in storage zone **410**, data objects D4, D5 and D6 are stored in storage zone **420**, and data objects D7, D8 and D9 are stored in storage zone **430**. Should one of the storage nodes in one of the storage zones become impaired, become unavailable, become inaccessible or go down, the system may recreate the data by using available data objects along with parity objects. When the data objects for the data item are distributed among the three zones, the system creates six parity objects P1 through P6. The system also creates reassembly instructions. The reassembly instructions, shown as R in a triangle in FIG. 4, provide information about how to recreate, reassemble or reconstitute the data item from (in this example, any nine of) the constituent data objects and parity objects. The reassembly instructions include which combinations of objects are needed to reassemble the data item. The reassembly instructions take up very little storage space. The same reassembly instructions are included in each zone. In this example, there are nine data objects and six parity objects. To recreate the data, a combination of any nine objects is needed. That is, for example, to recreate the data item, six data objects and three parity objects may be used. This would be the situation if one of the three storage zones shown in FIG. 4 cannot be used or accessed to recreate the data item. That is, for example, if zone **420** is down or inaccessible, the data objects D1, D2, D3, D6, D7, D8 and parity objects P1, P2, P5 from zones **410** and **430** may be used to recreate the data item.

The creation of parity objects may be based upon a number of well-known erasure coding techniques. In one embodiment, the parity objects are created with the Reed-Solomon technique. See James S. Plank "Erasure Codes for Storage Applications," Tutorial, FAST-2005: 4th Usenix Conference on File and Storage Technologies San Francisco, Calif., December, 2005 and James S. Plank, "A Tutorial on Reed-Solomon Coding for Fault-Tolerance in RAID-like Systems," Software—Practice & Experience, 27(9), September, 1997, pp. 995-1012.

The property of interest for such erasure coding techniques is the ability to repair the loss of data objects with an equal number of parity objects using the technique. For example, the calculation of parity can be used as a simple erasure code. Consider the 8 bit string "10011100", with the addition of the parity bit of "1" (even parity). If the system presents the string "1001X100" with the parity bit of "1" and the parity coding is known to be even, the receiver can infer that the missing bit "X" is a "1". This technique can be generalized to correct as many errors in the stored data as there are parity. The system and methods described herein are dependent on this general property of parity erasure codes and not a specific algorithm or implementation.

Referring to the examples shown in FIGS. 3 and 4, as the data item was partitioned into nine objects, in the fully replicated storage system 300 of FIG. 3, the amount of storage space required for a three zone system would be three times the size of the data. According to the resilient distributed replicated storage system described herein, a zone would have enough storage space for five objects and a small amount of storage space for reassembly instructions. The example resilient distributed replicated storage system 400 of FIG. 4 requires less than two-thirds of the capacity of the fully replicated storage system 300 of FIG. 5 while providing for resiliency and access to data. The one-third reduction in storage needed may result in significant monetary cost savings while providing the user a desired level of resiliency and accessibility.

The software running on a controller or computing device in storage zone 410 may monitor the health of the other storage zones 420 and 430 and/or the storage media in the storage zones 420 and 430. When an impending or actual problem or failure is detected in the storage zones 420 or 430, the storage zone 410 may identify this situation as an access alert. After identifying an access alert, the system uses data objects and parity objects to recreate the data according to the reassembly instructions. The access alert may be identified by a node, and the node may use the reassembly instructions to reassemble the data item by sending requests for data objects and parity objects to other nodes in other zones, if necessary.

In this way a customer of the resilient distributed replicated data storage system may receive the reliability or service level desired or required from a fully replicated storage system at a much lower cost.

The example shown in and described regarding FIG. 4 involves three storage zones and a single data item partitioned into nine objects. The techniques described herein apply to storage systems with multiple storage zones, regardless of the number of storage zones. The techniques described herein apply to any sized data item. The techniques described herein apply to systems with multiple data items. The techniques described herein require that the total number of data objects and parity objects can be divided evenly among the zones in the particular configuration of the resilient distributed replicated data storage system. The techniques described herein optionally allow for multiple copies of data objects and parity objects to be stored to meet the divided equally requirement among storage zones as well as to increase the resiliency of the stored object.

Description of Processes

The method used in the intelligent distributed replicated data storage system may be described as taking a data item, partitioning it into pieces known as data objects and calculating parity information known as parity objects along with reassembly instructions. The system allows for the spreading of data objects and parity objects among multiple storage zones to meet the availability and durability requirements of a particular application or customer. The spreading of the objects among zones is performed so that the durability of the data item is ensured. The number of data objects and parity objects is chosen to map onto the number of zones provided and the required durability. The system may duplicate parity and data objects to increase the resiliency of the stored item according to the policy requirements.

Referring now to FIG. 5, a flow chart of the actions taken to store or put data in a resilient distributed replicated data storage system is shown. To store or put a data item, the system receives a storage request that includes a data item (or reference or link to a data item) and may include storage criteria, as shown in block 510. The storage criteria may, in

different embodiments specify the storage policy for the data item by a code or name, or may include specific distribution instructions for the data item. The specific distribution instruction may include the exact number of data objects and parity objects the data item for the data item. In one implementation, the system predetermines the number of pieces the item data should be partitioned into according to information provided by the user of the system, and all data items are partitioned the same way. The number of pieces may be determined in part based on the storage criteria specified. The number of pieces is dependent on multiple factors, including the number of zones in the storage system, the specified or desired durability/resiliency of the system, and the desired speed of access to the data item. To increase the resiliency of the storage system, multiple copies of the data objects or parity objects may be dispersed among the storage zones. The number of parity objects may be determined in part based on the storage criteria specified. The number of parity objects is determined based on the number of pieces and the erasure coding technique used. The system may provide a user interface that allows the user to specify information for these factors by allowing for explicit entry of numbers (in text or numerical entry fields, radio button with associated text guides, and/or pull down menus), selection of factors by radio button, slider or other user interface item (such as, for example, a sliding scale for each of access speed, resiliency/durability). The number of pieces evaluation may be performed once during a system install, and may later be recalculated upon a system reconfiguration, such as adding or removing a zone, or changing of a desired access speed or resiliency. In another implementation, the number of pieces a data item is to be partitioned into may be hard coded or system defined.

The system then partitions the data item into data objects corresponding to the number of pieces, as shown in block 520. The system also creates parity objects for the data item, as shown in block 530. The partitioning into data objects and creating parity objects is achieved according to erasure coding techniques by one or more nodes in a zone in the distributed replicated data storage system.

According to erasure coding, the data item is divided into n pieces and recoded into $n+p$ total objects, namely n data objects and p parity objects, such that there will be t total objects. According to erasure coding, $t > n > p$. The key property of erasure coding is that the original item can be reconstructed from any n objects from the combination of data objects and parity objects, where the combined size for the n objects is equal or greater than the original data item size. For example, referring to FIG. 4, when the data item is partitioned into nine data objects n and four parity objects p , the total number of objects t is 10 and the number of objects needed to recreate the original data item is n or nine. The erasure coding used may be according to the Reed-Solomon techniques or other suitable techniques.

The system then creates reassembly instructions for the data item, as shown in block 540. The reassembly instructions include information about the data objects and the parity objects including their object identifiers that encode the location of the object.

The system then stores the data objects, parity objects and reassembly instructions in the current storage zone, as shown in block 550. The system stores the data objects, parity objects and reassembly instructions among the other storage zones according to the policy for the data item, as shown in block 560. In one embodiment the system performs the actions in blocks 550 and 560 concurrently when synchronized replication is required. The system performs the actions

in blocks 550 and 560 sequentially when asynchronous replication is sufficient. The system stores the data objects, parity objects and reassembly instructions among the storage zones so that when one zone goes down, the system can still provide the data item when requested. In one implementation, the data objects are spread evenly among the storage zones, and there is at least one parity object per storage zone. A copy of the reassembly instructions is stored in every storage zone; the reassembly instructions are the same in each storage zone.

In another version of the system, actions may be taken to avoid blocking, delays or other problems stemming from a disconnected, unavailable or otherwise inaccessible zone. For example, referring to FIG. 4, if storage zone 420 is inaccessible, rather than blocking or waiting for storage zone 420 to become available, the system would create a virtual copy of the objects and other pertinent information intended for the inaccessible storage zone (420) as a virtual zone at another storage zone, for example, in storage zone 410. FIG. 8 is a flow chart of additional actions taken to store or put data in a resilient distributed replicated data storage system. FIG. 9 is a flow chart of further additional actions taken to store or put data in a resilient distributed replicated data storage system, and is a continuation of FIG. 8.

Referring to FIG. 8, after or as part of initiating storing the data objects, parity objects and reassembly instructions among storage zones according to the policy for the data item as shown in block 560 of FIG. 5, the system values the latency for each storage zone to which the data item is to be stored, as shown in block 812. As used herein latency means the amount of time to transmit or transfer data from a first location to a second location, namely from a current zone to another zone. In some embodiments, a round trip latency may be used. For each storage zone to which the data item is to be stored, starting with the lowest latency zone, the system evaluates whether the zone is accessible, as shown in blocks 812 and 816. This may be done by pinging the zone or otherwise communicating with the zone. When the zone is accessible, the system stores the data objects, parity objects and reassembly instructions in the storage zones according to the policy for the data item, as shown in block 820.

When the zone is inaccessible, the system creates a virtual version of the inaccessible zone as a virtual zone, as shown in block 830. The system stores the data objects, parity objects and reassembly instructions according to the policy for the data item to the virtual zone as if it were the real zone to which the data item is to be stored, as shown in block 832. This continues for each zone in lowest latency order. Specifically, a check is made to determine whether there are more zones to which the data is to be stored, as shown in block 840. If there are more zones, the next (lowest latency) zone is selected, as shown in block 842, and the flow of actions continues at block 816. If storage has been completed to all zones (both virtual and real), the flow of actions continues at block 910 of FIG. 9.

After the data item has been stored to real and virtual zones pursuant to the policy for the data item, any objects stored in virtual zones need to be moved to and stored to the intended destination zone according to the policy for the data item. This may be achieved according to the actions described regarding FIG. 9. Referring now to FIG. 9, continuing from block 840 of FIG. 8, for each virtual zone to which the data item has been temporarily stored, a check is made to learn whether the zone associated with the virtual zone is accessible, as shown in blocks 910 and 912. This may be done by pinging the zone or otherwise communicating with the zone. If the real zone is accessible, the system moves the data objects, parity objects and reassembly instructions from the virtual zone to the intended destination zone, as shown in

block 920. If the real zone remain inaccessible, as shown in block 912, the system wait a system defined amount of time before checking for that zone again, as shown in block 930. After blocks 920 and 930, the system checks whether there are any further virtual zones (including the just checked zone), as shown in block 940. If there are virtual zones still in use, the next (or same when only one) virtual zone is selected, as shown in block 924, and the flow of actions continues at block 912 where a check is made to learn whether the real zone corresponding to the currently selected virtual zone is accessible. If there are no virtual zones in use, as shown in block 940, the flow of actions ceases.

Referring now to FIG. 6, a flow chart of the actions taken to read or get data included in a resilient distributed replicated data storage system is shown. The system receives a read request for a data item, as shown in block 610. The system determines the location of data objects and parity objects at lowest latency nodes needed to construct the data item by referring to the reassembly information, as shown in block 612. That is, the system retrieves and reviews the reassembly information to locate sufficient objects to reconstruct the requested data item by referring to a latency map. The system attempts to reconstruct the data item as quickly as possible.

Each node as it operates tracks the time it takes for operations with other nodes. This information about timing concerning operations with other nodes is stored in a list of nodes that includes the historical time to reach other nodes or perform operations involving other nodes. This list of nodes and times may be referred to as a latency map or map of latencies.

When attempting to reconstruct the data item as quickly as possible, the system may refer to the map of latencies to other nodes in the system. The object identifier includes the nodes where each object may be located. The system uses these two pieces of information, the node and the latency to the node, to attempt to retrieve sufficient objects to reconstruct the item in the shortest time. The system then attempts to obtain the data objects and parity objects for the data item, as shown in block 614.

If one or more lowest latency nodes at which data objects or parity objects are stored are either impaired or inaccessible as shown in block 616, the system determines the location of data objects and parity objects at not impaired and accessible nodes needed to construct the data item by referring to the reassembly information and obtains the objects, as shown in block 630. These may be referred to as alternate nodes, that is, the nodes that are not lowest latency but that are not impaired and are available. The system constructs the data item from data objects and parity objects obtained from not impaired and accessible nodes identified in the reassembly information, as shown in block 632, that is, the alternate nodes. The system then constructs the data item from data objects and parity objects obtained from nodes identified in the reassembly information, as shown in block 632. The system provides the requested data item to the requester, as shown in block 650.

When the system attempts to obtain the data objects and parity objects for the data item and all the nodes at which the data objects are stored are not impaired and accessible, the system obtains the data objects from lowest latency nodes identified in the reassembly information, as shown in block 620. The system then constructs the data item from data objects and parity objects obtained from nodes identified in the reassembly information, as shown in block 622. The system provides the requested data item to the requester, as shown in block 640.

When the system includes the use of virtual zones when intended storage zones are inaccessible, the actions described

11

in FIG. 6 may be augmented to include the actions described in FIG. 10. Referring to FIG. 10, after block 612 of FIG. 6, for each storage zone in which nodes storing data objects and parity objects are located, a check is made to learn whether the zone is accessible, as shown in block 1012. If the zone is accessible, the flow of actions returns to block 614 in FIG. 4, described above. If the zone is inaccessible, a check is made to learn whether there is a virtual zone available for the inaccessible zone, as shown in block 1020. If there is no virtual zone available, the flow of actions resumes at block 630 of FIG. 6, described above. If a virtual zone is available for the inaccessible zone, the system proceeds with using nodes in the virtual zone in place of the actual zone, as shown in block 1022. The flow of actions continues at block 620 of FIG. 6.

Alternatively a check may be made whether any of the zones where the data is stored are stored as virtual zones, and if so, the system proceeds with using nodes in the virtual zone in place of the actual zone.

The methods described above and shown in FIGS. 5, 6, 8, 9 and 10 may be performed by and managed by a node included in a zone in a distributed replicated data storage system in which the independent nodes communicate with each other in a peer to peer manner. In another embodiment, the methods described above and shown in FIGS. 5, 6, 8, 9 and 10 may be performed by and managed by an independent computing device or controller in a storage zone in a resilient distributed replicated data storage system.

To make the system described herein more resilient and more responsive, a local parity object (or objects) may be created at each zone to assist in recreating any one object in the zone that may not be accessible. For example, if the data objects and parity objects from a zone are needed and one of the parity objects or data objects is on a node or drive that is down, the inaccessible object can be created locally using a locally stored parity object (or objects). This configuration adds another layer of resiliency and reliability to the system.

Referring now to FIG. 7, a block diagram of a second example resilient distributed replicated data storage system 700 is shown. In the resilient distributed replicated data storage system 700, a local parity object is created and stored in each zone so that if an object stored in that zone is not accessible, it may be recreated locally. The use of local parity objects (L1, L2 and L3 as shown) increases the resiliency and data throughput of the system by alleviating the need to access objects from other zones to recreate an inaccessible object. In this example, a data item is partitioned into nine pieces, data objects D1 through D9. Data objects D1, D2 and D3 are stored in storage zone 710, data objects D4, D5 and D6 are stored in storage zone 720, and data objects D7, D8 and D9 are stored in storage zone 730. Should one of the storage zones or a node within a zone become impaired, become unavailable, become inaccessible or go down, the system may recreate the data by using available data objects along with parity objects. When the data objects for the data item are distributed among the three zones, the system creates six parity objects P1 through P6. In this extra resilient configuration, the system also creates a local parity object for each zone. The local parity objects are created using Reed-Solomon erasure coding or any suitable erasure coding scheme. As shown, the system creates local parity object L1 for zone 710 based on the data objects D1, D2, D3 and parity objects P1, P2 included in the zone. As shown, the system creates local parity object L2 for zone 720 based on the data objects D4, D5, D6 and parity objects P3, P4 included in the zone. As

12

shown, the system creates local parity object L3 for zone 730 based on the data objects D7, D8, D9 and parity objects P5, P6 included in the zone.

The system also creates reassembly instructions. The reassembly instructions, shown as R in a triangle in FIG. 7, provide information about how to recreate, reassemble or reconstitute the data item from (in this example, any nine of) the constituent data objects and parity objects. In this configuration, the reassembly instructions also include information needed to recreate an object inaccessible within a storage zone using a local parity object. The reassembly instructions take up very little space and are small in size. In one embodiment, the same reassembly instructions are included in each zone, and a separate local reassembly instruction object is also provided. In another embodiment, the system reassembly instructions are augmented with local reassembly instructions such that the reassembly instruction object is different for each zone.

When local parity objects are used, the flow of actions to store a data item shown in FIG. 5 is augmented by adding creating a local parity object or local parity objects between blocks 540 and 550. In addition, the creating reassembly instructions for the data item in block 550 may include creating assembly instructions for objects in each of the zones based on the local parity object or objects.

When local parity objects are used, the flow of actions to read or get a data item shown in FIG. 6 is augmented by adding identifying an inaccessible or corrupted drive or object in a node in the zone, and locally recreating the inaccessible object using the local parity object (or objects) according to local zone reassembly instructions. These actions may be taken between blocks 614 and 616, such that the node determines all nodes are not impaired. In this way, a local reconstituted version of the object is obtained from a local zone without having to seek the object from a node in a different zone. This results in increased throughput and efficiency of the system and adds a level of resiliency to the system.

Closing Comments

Throughout this description, the embodiments and examples shown should be considered as exemplars, rather than limitations on the apparatus and procedures disclosed or claimed. Although many of the examples presented herein involve specific combinations of method acts or system elements, it should be understood that those acts and those elements may be combined in other ways to accomplish the same objectives. With regard to flowcharts, additional and fewer steps may be taken, and the steps as shown may be combined or further refined to achieve the methods described herein. Acts, elements and features discussed only in connection with one embodiment are not intended to be excluded from a similar role in other embodiments.

As used herein, "plurality" means two or more.

As used herein, a "set" of items may include one or more of such items.

As used herein, whether in the written description or the claims, the terms "comprising", "including", "carrying", "having", "containing", "involving", and the like are to be understood to be open-ended, i.e., to mean including but not limited to. Only the transitional phrases "consisting of" and "consisting essentially of", respectively, are closed or semi-closed transitional phrases with respect to claims.

Use of ordinal terms such as "first", "second", "third", etc., "primary", "secondary", "tertiary", etc. in the claims to modify a claim element does not by itself connote any priority, precedence, or order of one claim element over another or the temporal order in which acts of a method are performed,

13

but are used merely as labels to distinguish one claim element having a certain name from another element having a same name (but for use of the ordinal term) to distinguish the claim elements.

As used herein, “and/or” means that the listed items are alternatives, but the alternatives also include any combination of the listed items.

It is claimed:

1. A system comprising:

a plurality of storage zones, each storage zone comprising a plurality of nodes wherein each node comprises a plurality of storage devices and a controller, the controller including a processor and memory

a first node of a plurality of nodes included in a first zone of the plurality of zones, the first node having instructions stored thereon which when executed cause a first processor included in a first controller in the first node to perform actions including:

receiving a storage request from a client of the system to store a data item;

partitioning the data item into a plurality of data objects, each object including a storage policy identifier and a data portion;

creating a plurality of parity objects for the data item;

creating reassembly instructions for the data item;

storing groups of some of the data objects and groups of some of the parity objects and the reassembly instructions among the nodes included in the storage zones according to a predetermined set of criteria, the storing including

evaluating a latency for each storage zone to which the data item is to be stored,

performing the storing in lowest latency order,

evaluating each storage zone to determine whether the storage zone is accessible,

when one of the storage zones is not accessible, creating a virtual zone in the first zone to store the data objects, parity objects and reassembly instructions intended for the inaccessible storage zone;

receiving a read request for the data item;

obtaining data objects and parity objects to reconstruct the data item from lowest latency nodes, including

when one of the storage zones having storage nodes where the data objects and parity objects to reconstruct the data item are stored are inaccessible, using the virtual zone in place of the inaccessible zone to obtain the data objects and parity objects;

providing the requested data item to a requester.

2. The system of claim 1 wherein the storing further comprises:

checking whether the storage zone corresponding to the virtual zone is accessible

when the storage zone corresponding to the virtual zone remains inaccessible, waiting a system defined amount of time

repeating the checking and the waiting until the storage zone corresponding to the virtual zone becomes accessible

when the storage zone corresponding to the virtual zone is accessible, moving the data objects, parity objects and reassembly instructions from the virtual zone to the intended storage zone.

3. The system of claim 1 wherein the storage devices are selected from the group including hard disk drives and solid state storage devices.

14

4. The system of claim 1 wherein the partitioning the data item partitions the data item into data objects wherein some of the data objects are duplicated among the storage zones.

5. The system of claim 1 wherein the partitioning the data item partitions the data object into data items that are equally divided among the storage zones.

6. The system of claim 1 wherein the partitioning the data item partitions the data object into data items that are unequally divided among the storage zones.

7. The system of claim 1 wherein the creating the parity objects creates parity objects wherein some of the parity objects are duplicated among the storage zones.

8. The system of claim 1 wherein the creating the parity objects creates parity objects that are equally divided among the storage zones.

9. The system of claim 1 wherein the creating the parity objects creates parity objects that are unequally divided among the storage zones.

10. The system of claim 1 wherein the obtaining data objects and parity objects to reconstruct the data item from lowest latency nodes further includes

checking whether any storage zones having storage nodes where the data objects and parity objects to reconstruct the data item are stored are inaccessible.

11. The system of claim 1 wherein obtaining data objects and parity objects to reconstruct the data item from lowest latency nodes further includes:

when the lowest latency nodes are not impaired and accessible, constructing the data item from the data objects and parity object at the lowest latency nodes according to the reassembly instructions

when one of the lowest latency nodes is either impaired or not accessible, constructing the data from the data objects and parity objects obtained from alternative nodes according to the reassembly instructions

using a local parity object to allow for recreation of an inaccessible or corrupted one of the parity objects in the group of parity objects or one of the data objects in the group of data objects.

12. The system of claim 1 wherein when one of the lowest latency nodes is either impaired or not accessible, reconstituting the node that is either impaired or not accessible using a local parity object according to local reassembly instructions included in the reassembly instructions.

13. A system comprising:

a plurality of storage zones, each storage zone comprising a plurality of storage nodes wherein each storage node comprises a plurality of storage devices and a controller, the controller including a processor and memory

a first node of a plurality of nodes included in a first zone of the plurality of zones, the first node having instructions stored thereon which when executed cause a first processor included in a first controller in the first node to perform actions including:

receiving a storage request from a client of the system to store a data item, the storage request including (a) the data item or a data item reference and (b) a storage criteria;

partitioning the data item into a plurality of data objects, each object including a storage policy identifier and a data portion according to the storage criteria;

creating a plurality of parity objects for the data item according to the storage criteria;

creating reassembly instructions for the data item according to the storage criteria;

storing the data objects and the parity objects for the data item among the nodes among the storage zones

15

according to the storage criteria and storing the reassembly instructions among the storage nodes among the storage zones, the storing including
 evaluating a latency for each storage zone to which the data item is to be stored,
 performing the storing in lowest latency order,
 evaluating each storage zone to determine whether the storage zone is accessible,
 when one of the storage zones is not accessible, creating a virtual zone in the first zone to store the data objects, parity objects and reassembly instructions intended for the inaccessible storage zone;
 receiving a read request for the data item;
 obtaining data objects and parity objects to reconstruct the data item from lowest latency nodes, including
 when one of the storage zones having storage nodes where the data objects and parity objects to reconstruct the data item are stored are inaccessible, using the virtual zone in place of the inaccessible zone to obtain the data objects and parity objects;
 providing the requested data item to a requester.

14. The system of claim 13 wherein the storing further comprises:
 checking whether the storage zone corresponding to the virtual zone is accessible
 when the storage zone corresponding to the virtual zone remains inaccessible, waiting a system defined amount of time
 repeating the checking and the waiting until the storage zone corresponding to the virtual zone becomes accessible
 when the storage zone corresponding to the virtual zone is accessible, moving the data objects, parity objects and reassembly instructions from the virtual zone to the intended storage zone.

15. The system of claim 13 wherein the storage devices are selected from the group including hard disk drives and solid state storage devices.

16. The system of claim 13 wherein the storage criteria specify whether the data objects and the parity objects should

16

be equally divided among the storage zones, unequally divided among the storage zones, and duplicated among the storage zones.

17. The system of claim 13 wherein the storage criteria specify redundancy, latency, and/or resiliency characteristics.

18. The system of claim 13 wherein the storage criteria include designated placement for the data objects and the parity objects among the storage zones.

19. The system of claim 13 wherein the storage criteria are designations of redundancy, latency, and/or resiliency characteristics that are processed to determine the distribution of the data objects and the parity objects among the storage zones.

20. The system of claim 13 wherein obtaining data objects and parity objects to reconstruct the data item from lowest latency nodes includes

checking whether any storage zones having storage nodes where the data objects and parity objects to reconstruct the data item are stored are inaccessible.

21. The system of claim 13 wherein the obtaining further comprises:

when the lowest latency nodes are not impaired and accessible, constructing the data item from the data objects and parity object at the lowest latency nodes according to the reassembly instructions

when one of the lowest latency nodes is either impaired or not accessible, constructing the data from the data objects and parity objects obtained from alternative nodes according to the reassembly instructions

using a local parity object to allow for recreation of an inaccessible or corrupted one of the parity objects in the group of parity objects or one of the data objects in the group of data objects.

22. The system of claim 13 wherein when one of the lowest latency nodes is either impaired or not accessible, reconstituting the node that is either impaired or not accessible using a local parity object according to local reassembly instructions included in the reassembly instructions.

* * * * *